arXiv:1607.00847v8 [cs.LG] 24 Nov 2016

Confidence-Weighted Bipartite Ranking

Majdi Khalid, Indrakshi Ray, and Hamidreza Chitsaz

Colorado State University, Fort Collins, USA
majdi.khaled@gmail.com, indrakshi.ray@colostate.edu,
chitsaz@chitsazlab.org
www.colostate.edu

Abstract. Bipartite ranking is a fundamental machine learning and data mining problem. It commonly concerns the maximization of the AUC metric. Recently, a number of studies have proposed online bipartite ranking algorithms to learn from massive streams of class-imbalanced data. These methods suggest both linear and kernel-based bipartite ranking algorithms based on first and second-order online learning. Unlike kernelized ranker, linear ranker is more scalable learning algorithm. The existing linear online bipartite ranking algorithms lack either handling non-separable data or constructing adaptive large margin. These limitations yield unreliable bipartite ranking performance. In this work, we propose a linear online confidence-weighted bipartite ranking algorithm (CBR) that adopts soft confidence-weighted learning. The proposed algorithm leverages the same properties of soft confidence-weighted learning in a framework for bipartite ranking. We also develop a diagonal variation of the proposed confidence-weighted bipartite ranking algorithm to deal with high-dimensional data by maintaining only the diagonal elements of the covariance matrix. We empirically evaluate the effectiveness of the proposed algorithms on several benchmark and high-dimensional datasets. The experimental results validate the reliability of the proposed algorithms. The results also show that our algorithms outperform or are at least comparable to the competing online AUC maximization methods.

Keywords: online ranking, imbalanced learning, AUC maximization

1 Introduction

Bipartite ranking is a fundamental machine learning and data mining problem because of its wide range of applications such as recommender systems, information retrieval, and bioinformatics [1,25,29]. Bipartite ranking has also been shown to be an appropriate learning algorithm for imbalanced data [6]. The aim of the bipartite ranking algorithm is to maximize the Area Under the Curve (AUC) by learning a function that scores positive instances higher than negative instances. Therefore, the optimization problem of such a ranking model is formulated as the minimization of a pairwise loss function. This ranking problem can be solved by applying a binary classifier to pairs of positive and negative instances, where

the classification function learns to classify a pair as positive or negative based on the first instance in the pair. The key problem of this approach is the high complexity of a learning algorithm that grows quadratically or subquadratically [19] with respect to the number of instances.

Recently, significant efforts have been devoted to developing scalable bipartite ranking algorithms to optimize AUC in both batch and online settings [12,15,17,24,37]. Online learning approach is an appealing statistical method because of its scalability and effectivity. The online bipartite ranking algorithms can be classified on the basis of the learned function into linear and nonlinear ranking models. While they are advantageous over linear online ranker in modeling nonlinearity of the data, kernelized online ranking algorithms require a kernel computation for each new training instance. Further, the decision function of the nonlinear kernel ranker depends on support vectors to construct the kernel and to make a decision.

Online bipartite ranking algorithms can also be grouped into two different schemes. The first scheme maintains random instances from each class label in a finite buffer [20,37]. Once a new instance is received, the buffer is updated based on stream oblivious policies such as Reservoir Sampling (RS) and First-In-First-Out (FIFO). Then the ranker function is updated based on a pair of instances, the new instance and each opposite instance stored in the corresponding buffer. These online algorithms are able to deal with non-separable data, but they are based on simple first-order learning. The second approach maintains the first and second statistics [15], and is able to adapt the ranker to the importance of the features [12]. However, these algorithms assume the data are linearly separable.

Moreover, these algorithms make no attempt to exploit the confidence information, which has shown to be very effective in ameliorating the classification performance [8,13,35]. Confidence-weighted (CW) learning takes the advantage of the underlying structure between features by modeling the classifier (i.e., the weight vector) as a Gaussian distribution parameterized by a mean vector and covariance matrix [13]. This model captures the notion of confidence for each weight coordinate via the covariance matrix. A large diagonal value corresponding to the i-th feature in the covariance matrix results in less confidence in its weight (i.e., its mean). Therefore, an aggressive update is performed on the less confident weight coordinates. This is analogous to the adaptive subgradient method [14] that involves the geometric structure of the data seen so far in regularizing the weights of sparse features (i.e., less occurring features) as they are deemed more informative than dense features. The confidence-weighted algorithm [13] has also been improved by introducing the adaptive regularization (AROW) that deals with inseparable data [10]. The soft confidence-weighted (SCW) algorithm improves upon AROW by maintaining an adaptive margin [35].

In this paper, we propose a novel framework that solves a linear online bipartite ranking using the soft confidence-weighted algorithm [35]. The proposed confidence-weighted bipartite ranking algorithm (CBR) entertains the fast training and testing phases of linear bipartite ranking. It also enjoys the capabil-

ity of the soft confidence-weighted algorithm in learning confidence-weighted model, handling linearly inseparable data, and constructing adaptive large margin. The proposed framework follows an online bipartite ranking scheme that maintains a finite buffer for each class label while updating the buffer by one of the stream oblivious policies such as Reservoir Sampling (RS) and First-In-First-Out (FIFO) [32]. We also provide a diagonal variation (CBR-diag) of the proposed algorithm to handle high-dimensional datasets.

The remainder of the paper is organized as follows. In Section 2, we briefly review closely related work. We present the confidence-weighted bipartite ranking (CBR) and its diagonal variation (CBR-diag) in Section 3. The experimental results are presented in Section 4. Section 5 concludes the paper and presents some future work.

2 Related Work

The proposed bipartite ranking algorithm is closely related to the online learning and bipartite ranking algorithms. What follows is a brief review of recent studies related to these topics.

Online Learning. The proliferation of big data and massive streams of data emphasize the importance of online learning algorithms. Online learning algorithms have shown a comparable classification performance compared to batch learning, while being more scalable. Some online learning algorithms, such as the Perceptron algorithm [30], the Passive-Aggressive (PA) [7], and the online gradient descent [38], update the model based on a first-order learning approach. These methods do not take into account the underlying structure of the data during learning. This limitation is addressed by exploring second-order information to exploit the underlying structure between features in ameliorating learning performance [4,10,13,14,27,35]. Moreover, kernelized online learning methods have been proposed to deal with nonlinearly distributed data [3,11,28].

Bipartite Ranking. Bipartite ranking learns a real-valued function that induces an order on the data in which the positive instances precede negative instances. The common measure used to evaluate the success of the bipartite ranking algorithm is the AUC [16]. Hence the minimization of the bipartite ranking loss function is equivalent to the maximization of the AUC metric. The AUC presents the probability that a model will rank a randomly drawn positive instance higher than a random negative instance. In batch setting, a considerable amount of studies have investigated the optimization of linear and nonlinear kernel ranking function [5,18,22,23]. More scalable methods are based on stochastic or online learning [31,33,34]. However, these methods are not specifically designed to optimize the AUC metric. Recently, a few studies have focused on the optimization of the AUC metric in online setting. The first approach adopted a framework that maintained a buffer with limited capacity to store random instances to deal with the pairwise loss function [20,37]. The other methods maintained only the first and second statistics for each received instance and optimized the AUC in one pass over the training data [12,15]. The work [12]

exploited the second-order technique [14] to make the ranker aware of the importance of less frequently occurring features, hence updating their weights with a higher learning rate.

Our proposed method follows the online framework that maintains fixed-sized buffers to store instances from each class label. Further, it exploits the online second-order method [35] to learn a robust bipartite ranking function. This distinguishes the proposed method from [20,37] employing first-order online learning. Also, the proposed method is different from [12,15] by learning a confidence-weighted ranker capable of dealing with non-separable data and learning an adaptive large margin. The most similar approaches to our method are [33,34]. However, these are not designed to directly maximize the AUC. They also use classical first-order and second-order online learning whereas we use the soft variation of confidence-weighted learning that has shown a robust performance in the classification task [35].

3 Online Confidence-Weighted Bipartite Ranking

3.1 Problem Setting

We consider a linear online bipartite ranking function that learns on imbalanced data to optimize the AUC metric [16]. Let $\mathcal{S} = \{x_i^+ \cup x_j^- \in \mathcal{R}^d | i = \{1,\ldots,n\}, j = \{1,\ldots,m\}\}$ denotes the input space of dimension d generated from unknown distribution $\mathcal{D}$, where x_i^+ is the i-th positive instance and x_j^- is the j-th negative instance. The n and m denote the number of positive and negative instances, respectively. The linear bipartite ranking function $f : \mathcal{S} \to \mathcal{R}$ is a real valued function that maximizes the AUC metric by minimizing the following loss function:

$$\mathcal{L}(f;\mathcal{S}) = \frac{1}{nm}\sum_{i=1}^{n}\sum_{j=1}^{m} I(f(x_i^+) \leq f(x_j^-)),$$

where $f(x) = w^T x$ and $I(\cdot)$ is the indicator function that outputs 1 if the condition is held, and 0 otherwise. It is common to replace the indicator function with a convex surrogate function,

$$\mathcal{L}(f;\mathcal{S}) = \frac{1}{nm}\sum_{i=1}^{n}\sum_{j=1}^{m} \ell(f(x_i^+) - f(x_j^-)),$$

where $\ell(z) = max(0, 1-z)$ is the hinge loss function. Therefore, the objective function we seek to minimize with respect to the linear function is defined as follows:

$$\frac{1}{2}||w||^2 + C\sum_{i=1}^{n}\sum_{j=1}^{m} max(0, 1 - w(x_i^+ - x_j^-)), \tag{1}$$

where $||w||$ is the euclidean norm to regularize the ranker, and C is the hyper-parameter to penalize the sum of errors. It is easy to see that the optimization problem (1) will grow quadratically with respect to the number of training instances. Following the approach suggested by [37] to deal with the complexity of the pairwise loss function, we reformulate the objective function (1) as a sum of two losses for a pair of instances,

$$\frac{1}{2}||w||^2 + C\sum_{t=1}^{T} I_{(y_t=+1)} g_t^+(w) + I_{(y_t=-1)} g_t^-(w), \tag{2}$$

where $T = n + m$, and $g_t(w)$ is defined as follows

$$g_t^+(w) = \sum_{t'=1}^{t-1} I_{(y_{t'}=-1)} max(0, 1 - w(x_t - x_{t'})), \tag{3}$$

$$g_t^-(w) = \sum_{t'=1}^{t-1} I_{(y_{t'}=+1)} max(0, 1 - w(x_{t'} - x_t)). \tag{4}$$

Instead of maintaining all the received instances to compute the gradients $\nabla g_t(w)$, we store random instances from each class in the corresponding buffer. Therefore, two buffers B_+ and B_- with predefined capacity are maintained for positive and negative classes, respectively. The buffers are updated using a stream oblivious policy. The current stored instances in a buffer are used to update the classifier as in equation (2) whenever a new instance from the opposite class label is received.

The framework of the online confidence-weighted bipartite ranking is shown in Algorithm 1. The two main components of this framework are UpdateBuffer and UpdateRanker, which are explained in the following subsections.

3.2 Update Buffer

One effective approach to deal with pairwise learning algorithms is to maintain a buffer with a fixed capacity. This raises the problem of updating the buffer to store the most informative instances. In our online Bipartite ranking framework, we investigate the following two stream oblivious policies to update the buffer:

Reservoir Sampling **(RS)**: Reservoir Sampling is a common oblivious policy to deal with streaming data [32]. In this approach, the new instance (x_t, y_t) is added to the corresponding buffer if its capacity is not reached, $|B_{y_t}^t| < M_{y_t}$. If the buffer is at capacity, it will be updated with probability $\frac{M_{y_t}}{M_{y_t}^{t+1}}$ by randomly replacing one instance in $B_{y_t}^t$ with x_t.

First-In-First-Out **(FIFO)**: This simple strategy replaces the oldest instance with the new instance if the corresponding buffer reaches its capacity. Otherwise, the new instance is simply added to the buffer.

Algorithm 1: A Framework for Confidence-Weighted Bipartite Ranking (CBR)

Input:

- the penalty parameter C
- the capacity of the buffers M_+ and M_-
- η parameter
- $a_i = 1$ for $i \in 1, \ldots, d$

Initialize: $\mu_1 = \{0, \ldots, 0\}^d$, $B_+ = B_- = \emptyset$, $M_+^1 = M_-^1 = 0$
$\Sigma_1 = diag(a)$ or $G_1 = a$
for t = 1, . . . , T **do**
Receive a training instance (x_t, y_t)
if $y_t = +1$ **then**
$M_+^{t+1} = M_+^t + 1$, $M_-^{t+1} = M_-^t$, $B_-^{t+1} = B_-^t$
$B_+^{t+1} = \text{UpdateBuffer}(x_t, B_+^t, M_+^{t+1}, M_+)$
$[\mu_{t+1}, \Sigma_{t+1}] = \text{UpdateRanker}(\mu_t, \Sigma_t, x_t, y_t, C, B_-^{t+1}, \eta)$ or
$[\mu_{t+1}, G_{t+1}] = \text{UpdateRanker}(\mu_t, G_t, x_t, y_t, C, B_-^{t+1}, \eta)$ (CBR-diag)
else
$M_-^{t+1} = M_-^t + 1$, $M_+^{t+1} = M_+^t$, $B_+^{t+1} = B_+^t$
$B_-^{t+1} = \text{UpdateBuffer}(x_t, B_-^t, M_-^{t+1}, M_-)$
$[\mu_{t+1}, \Sigma_{t+1}] = \text{UpdateRanker}(\mu_t, \Sigma_t, x_t, y_t, C, B_+^{t+1}, \eta)$ or
$[\mu_{t+1}, G_{t+1}] = \text{UpdateRanker}(\mu_t, G_t, x_t, y_t, C, B_+^{t+1}, \eta)$ (CBR-diag)
end if
end for

3.3 Update Ranker

Inspired by the robust performance of second-order learning algorithms, we apply the soft confidence-weighted learning approach [35] to updated the bipartite ranking function. Therefore, our confidence-weighted bipartite ranking model (CBR) is formulated as a ranker with a Gaussian distribution parameterized by mean vector $\mu \in \mathcal{R}^d$ and covariance matrix $\Sigma \in \mathcal{R}^{d \times d}$. The mean vector μ represents the model of the bipartite ranking function, while the covariance matrix captures the confidence in the model. The ranker is more confident about the model value μ_p as its diagonal value $\Sigma_{p,p}$ is smaller. The model distribution is updated once the new instance is received while being close to the old model distribution. This optimization problem is performed by minimizing the Kullback-Leibler divergence between the new and the old distributions of the model. The online confidence-weighted bipartite ranking (CBR) is formulated as follows:

$$(\mu_{t+1}, \Sigma_{t+1}) = \underset{\mu, \Sigma}{\text{argmin}} D_{KL}(\mathcal{N}(\mu, \Sigma) || \mathcal{N}(\mu_t, \Sigma_t)) + C\ell^{\phi}(\mathcal{N}(\mu, \Sigma); (z, y_t)), \tag{5}$$

where $z = (x_t - x)$, C is the the penalty hyperparamter, $\phi = \Phi^{-1}(\eta)$, and Φ is the normal cumulative distribution function. The loss function $\ell^{\phi}(\cdot)$ is defined as:

$$\ell^{\phi}(\mathcal{N}(\mu, \Sigma); (z, y_t)) = max(0, \phi\sqrt{z^T \Sigma z} - y_t \mu \cdot z).$$

The solution of (5) is given by the following proposition.

Proposition 1. *The optimization problem (5) has a closed-form solution as follows:*

$$\mu_{t+1} = \mu_t + \alpha_t y_t \Sigma_t z,$$
$$\Sigma_{t+1} = \Sigma_t - \beta_t \Sigma_t z^T z \Sigma_t.$$

The coefficients α and β are defined as follows:

$\alpha_t = min\{C, max\{0, \frac{1}{\upsilon_t \zeta}(-m_t \psi + \sqrt{m_t^2 \frac{\phi^4}{4} + \upsilon_t \phi^2 \zeta})\}\}$,

$\beta_t = \frac{\alpha_t \phi}{\sqrt{u_t} + \upsilon_t \alpha_t \phi}$. *where* $u_t = \frac{1}{4}(-\alpha_t \upsilon_t \phi + \sqrt{\alpha_t^2 \upsilon_t^2 \phi^2 + 4\upsilon_t})^2$,
$\upsilon_t = z^T \Sigma_t z$, $m_t = y_t(\mu_t \cdot z)$, $\phi = \Phi^{-1}(\eta)$, $\psi = 1 + \frac{\phi^2}{2}$, $\zeta = 1 + \phi^2$, *and* $z = x_t - x$.

The proposition (1) is analogous to the one derived in [35].

Though modeling the full covariance matrix lends the CW algorithms a powerful capability in learning [9,26,35], it raises potential concerns with high-dimensional data. The covariance matrix grows quadratically with respect to the data dimension. This makes the CBR algorithm impractical with high-dimensional data due to high computational and memory requirements.

We remedy this deficiency by a diagonalization technique [9,14]. Therefore, we present a diagonal confidence-weighted bipartite ranking (CBR-diag) that models the ranker as a mean vector $\mu \in \mathcal{R}^d$ and diagonal matrix $\hat{\Sigma} \in \mathcal{R}^{d \times d}$. Let G denotes $diag(\hat{\Sigma})$, and the optimization problem of CBR-diag is formulated as follows:

$$(\mu_{t+1}, G_{t+1}) = \underset{\mu, G}{\operatorname{argmin}} D_{KL}(\mathcal{N}(\mu, G) || \mathcal{N}(\mu_t, G_t)) + C\ell^{\phi}(\mathcal{N}(\mu, G); (z, y_t)). \tag{6}$$

Proposition 2. *The optimization problem (6) has a closed-form solution as follows:*

$$\mu_{t+1} = \mu_t + \frac{\alpha_t y_t z}{G_t},$$
$$G_{t+1} = G_t + \beta_t z^2.$$

The coefficients α *and* β *are defined as follows*

$$\alpha_t = min\{C, max\{0, \frac{1}{\upsilon_t \zeta}(-m_t\psi + \sqrt{m_t^2 \frac{\phi^4}{4} + \upsilon_t \phi^2 \zeta})\}\},$$

$$\beta_t = \frac{\alpha_t \phi}{\sqrt{u_t} + \upsilon_t \alpha_t \phi},$$

where $u_t = \frac{1}{4}(-\alpha_t \upsilon_t \phi + \sqrt{\alpha_t^2 \upsilon_t^2 \phi^2 + 4\upsilon_t})^2$, $\upsilon_t = \sum_{i=1}^{d} \frac{z_i^2}{G_i + C}$, $m_t = y_t(\mu_t \cdot z)$, $\phi = \Phi^{-1}(\eta)$, $\psi = 1 + \frac{\phi^2}{2}$, $\zeta = 1 + \phi^2$, *and* $z = x_t - x$.

The propositions 1 and 2 can be proved similarly to the proof in [35]. The steps of updating the online confidence-weighted bipartite ranking with full covariance matrix or with the diagonal elements are summarized in Algorithm 2.

Algorithm 2: Update Ranker

Input:

- μ_t : current mean vector
- Σ_t *or* G_t : current covariance matrix or diagonal elements
- (x_t, y_t) : a training instance
- B : the buffer storing instances from the opposite class label
- C : penalty parameter
- η : the predefined probability

Output: updated ranker:

- μ_{t+1}
- Σ_{t+1} *or* G_{t+1}

Initialize: $\mu^1 = \mu_t$, ($\Sigma^1 = \Sigma_t$ or $G^1 = G_t$), $i = 1$
for $x \in B$ **do**
  Update the ranker (μ^i, Σ^i) with $z = x_t - x$ and y_t by
  $(\mu^{i+1}, \Sigma^{i+1}) = \underset{\mu,\Sigma}{\text{argmin}} D_{KL}(\mathcal{N}(\mu, \Sigma)||\mathcal{N}(\mu^i, \Sigma^i)) + C\ell^{\phi}(\mathcal{N}(\mu, \Sigma); (z, y_t))$
  or
  Update the ranker (μ^i, G^i) with $z = x_t - x$ and y_t by
  $(\mu^{i+1}, G^{i+1}) = \underset{\mu,G}{\text{argmin}} D_{KL}(\mathcal{N}(\mu, G)||\mathcal{N}(\mu^i, G^i)) + C\ell^{\phi}(\mathcal{N}(\mu, G); (z, y_t))$
  $i = i + 1$
end for
Return $\mu_{t+1} = \mu^{|B|+1}$
  $\Sigma_{t+1} = \Sigma^{|B|+1}$ or $G_{t+1} = G^{|B|+1}$

4 Experimental Results

In this section, we conduct extensive experiments on several real world datasets in order to demonstrate the effectiveness of the proposed algorithms. We also

compare the performance of our methods with existing online learning algorithms in terms of AUC and classification accuracy at the optimal operating point of the ROC curve (OPTROC). The running time comparison is also presented.

4.1 Real World Datasets

We conduct extensive experiments on various benchmark and high-dimensional datasets. All datasets can be downloaded from LibSVM[1] and the machine learning repository UCI[2] except the Reuters[3] dataset that is used in [2]. If the data are provided as training and test sets, we combine them together in one set. For cod-rna data, only the training and validation sets are grouped together. For rcv1 and news20, we only use their training sets in our experiments. The multi-class datasets are transformed randomly into class-imbalanced binary datasets. Tables 1 and 2 show the characteristics of the benchmark and the high-dimensional datasets, respectively.

Table 1. Benchmark datasets

Data	#inst	#feat	Data	#inst	#feat	Data	#inst	#feat
glass	214	10	cod-rna	331,152	8	australian	690	14
ionosphere	351	34	spambase	4,601	57	diabetes	768	8
german	1,000	24	covtype	581,012	54	acoustic	78,823	50
svmguide4	612	10	magic04	19,020	11	vehicle	846	18
svmguide3	1284	21	heart	270	13	segment	2,310	19

Table 2. High-dimensional datasets

Data	#inst	#feat
farm-ads	4,143	54,877
rcv1	15,564	47,236
sector	9,619	55,197
real-sim	72,309	20,958
news20	15,937	62,061
Reuters	8,293	18,933

4.2 Compared Methods and Model Selection

Online Uni-Exp [21]: An online pointwise ranking algorithm that optimizes the weighted univariate exponential loss. The learning rate is tuned by 2-fold cross validation on the training set by searching in $2^{[-10:10]}$.

[1] https://www.csie.ntu.edu.tw/ cjlin/libsvmtools/
[2] https://archive.ics.uci.edu/ml/
[3] http://www.cad.zju.edu.cn/home/dengcai/Data/TextData.html

OPAUC [15]: An online learning algorithm that optimizes the AUC in one-pass through square loss function. The learning rate is tuned by 2-fold cross validation by searching in $2^{[-10:10]}$, and the regularization hyperparameter is set to a small value 0.0001.
OPAUCr [15]: A variation of OPAUC that approximates the covariance matrices using low-rank matrices. The model selection step is carried out similarly to OPAUC, while the value of rank τ is set to 50 as suggested in [15].
$\mathbf{OAM_{seq}}$ [37]: The online AUC maximization (OAM) is the state-of-the-art first-order learning method. We implement the algorithm with the Reservoir Sampling as a buffer updating scheme. The size of the positive and negative buffers is fixed at 50. The penalty hyperparameter C is tuned by 2-fold cross validation on the training set by searching in $2^{[-10:10]}$.
AdaOAM [12]: This is a second-order AUC maximization method that adapts the classifier to the importance of features. The smooth hyperparameter δ is set to 0.5, and the regularization hyperparameter is set to 0.0001. The learning rate is tuned by 2-fold cross validation on the training set by searching in $2^{[-10:10]}$.
$\mathbf{CBR_{RS}}$ and **$\mathbf{CBR_{FIFO}}$**: The proposed confidence-weighted bipartite ranking algorithms with the Reservoir Sampling and First-In-First-Out buffer updating policies, respectively. The size of the positive and negative buffers is fixed at 50. The hyperparameter η is set to 0.7, and the penalty hyperparameter C is tuned by 2-fold cross validation by searching in $2^{[-10:10]}$.
CBR-diag$_{\mathbf{FIFO}}$: The proposed diagonal variation of confidence-weighted bipartite ranking that uses the First-In-First-Out policy to update the buffer. The buffers are set to 50, and the hyperparameters are tuned similarly to $\mathrm{CBR}_{\mathrm{FIFO}}$.

For a fair comparison, the datasets are scaled similarly in all experiments. We randomly divide each dataset into 5 folds, where 4 folds are used for training and one fold is used as a test set. For benchmark datasets, we randomly choose 8000 instances if the data exceeds this size. For high-dimensional datasets, we limit the sample size of the data to 2000 due to the high dimensionality of the data. The results on the benchmark and the high-dimensional datasets are averaged over 10 and 5 runs, respectively. A random permutation is performed on the datasets with each run. All experiments are conducted with Matlab 15 on a workstation computer with 8x2.6G CPU and 32 GB memory.

4.3 Results on Benchmark Datasets

The comparison in terms of AUC is shown in Table 3, while the comparison in terms of classification accuracy at OPTROC is shown in Table 4. The running time (in milliseconds) comparison is illustrated in Figure 1.

The results show the robust performance of the proposed methods CBR_{RS} and CBR_{FIFO} in terms of AUC and classification accuracy compared to other first and second-order online learning algorithms. We can observe that the improvement of the second-order methods such as OPAUC and AdaOAM over first-order method $\mathrm{OAM}_{\mathrm{seq}}$ is not reliable, while our CBR algorithms often outperform the $\mathrm{OAM}_{\mathrm{seq}}$. Also, the proposed methods are faster than $\mathrm{OAM}_{\mathrm{seq}}$, while

they incur more running time compared to AdaOAM except with spambase, covtype, and acoustic datasets. The pointwise method online Uni-Exp maintains fastest running time, but at the expense of the AUC and classification accuracy. We also notice that the performance of CBR_{FIFO} is often slightly better than CBR_{RS} in terms of AUC, classification accuracy, and running time.

Table 3. Comparison of AUC performance on benchmark datasets

Data	CBR_{RS}	CBR_{FIFO}	Online Uni-Exp	OPAUC	OAM_{seq}	AdaOAM
glass	**0.825** ± 0.043	0.823 ± 0.046	0.714 ± 0.075	0.798 ± 0.061	0.805 ± 0.047	0.794 ± 0.061
ionosphere	0.950 ± 0.027	**0.951** ± 0.028	0.913 ± 0.018	0.943 ± 0.026	0.946 ± 0.025	0.943 ± 0.029
german	**0.782** ± 0.024	0.780 ± 0.019	0.702 ± 0.032	0.736 ± 0.034	0.731 ± 0.028	0.770 ± 0.024
svmguide4	0.969 ± 0.013	**0.974** ± 0.013	0.609 ± 0.096	0.733 ± 0.056	0.771 ± 0.063	0.761 ± 0.053
svmguide3	0.755 ± 0.022	**0.764** ± 0.036	0.701 ± 0.025	0.737 ± 0.029	0.705 ± 0.033	0.738 ± 0.033
cod-rna	0.983 ± 0.000	**0.984** ± 0.000	0.928 ± 0.000	0.927 ± 0.001	0.951 ± 0.025	0.927 ± 0.000
spambase	0.941 ± 0.006	**0.942** ± 0.006	0.866 ± 0.016	0.849 ± 0.020	0.897 ± 0.043	0.862 ± 0.011
covtype	0.816 ± 0.003	**0.835** ± 0.001	0.705 ± 0.033	0.711 ± 0.041	0.737 ± 0.023	0.770 ± 0.010
magic04	0.799 ± 0.006	**0.801** ± 0.006	0.759 ± 0.006	0.748 ± 0.033	0.757 ± 0.015	0.773 ± 0.006
heart	0.908 ± 0.019	**0.909** ± 0.021	0.733 ± 0.039	0.788 ± 0.054	0.806 ± 0.059	0.799 ± 0.079
australian	0.883 ± 0.028	**0.889** ± 0.019	0.710 ± 0.130	0.735 ± 0.138	0.765 ± 0.107	0.801 ± 0.037
diabetes	0.700 ± 0.021	**0.707** ± 0.033	0.633 ± 0.036	0.667 ± 0.041	0.648 ± 0.040	0.675 ± 0.034
acoustic	0. 879 ± 0.006	**0.892** ± 0.003	0.876 ± 0.003	0.878 ± 0.003	0.863 ± 0.011	0.882 ± 0.003
vehicle	**0.846** ± 0.031	**0.846** ± 0.034	0.711 ± 0.053	0.764 ± 0.073	0.761 ± 0.078	0.792 ± 0.049
segment	0.900 ± 0.013	**0.903** ± 0.008	0.689 ± 0.061	0.828 ± 0.024	0.812 ± 0.035	0.855 ± 0.008

Table 4. Comparison of classification accuracy at OPTROC on benchmark datasets

Data	CBR_{RS}	CBR_{FIFO}	Online Uni-Exp	OPAUC	OAM_{seq}	AdaOAM
glass	**0.813** ± 0.044	0.811 ± 0.049	0.732 ± 0.060	0.795 ± 0.046	0.788 ± 0.040	0.783 ± 0.047
ionosphere	**0.946** ± 0.028	**0.946** ± 0.022	0.902 ± 0.028	0.936 ± 0.018	0.943 ± 0.017	0.938 ± 0.018
german	0.780 ± 0.022	**0.787** ± 0.019	0.741 ± 0.027	0.754 ± 0.022	0.751 ± 0.028	0.770 ± 0.030
svmguide4	0.951 ± 0.014	**0.956** ± 0.012	0.829 ± 0.021	0.843 ± 0.024	0.839 ± 0.022	0.848 ± 0.020
svmguide3	0.784 ± 0.015	**0.793** ± 0.016	0.784 ± 0.019	0.777 ± 0.024	0.780 ± 0.020	0.777 ± 0.024
cod-rna	0.948 ± 0.002	**0.949** ± 0.000	0.887 ± 0.001	0.887 ± 0.001	0.910 ± 0.019	0.887 ± 0.001
spambase	**0.899** ± 0.009	0.898 ± 0.009	0.818 ± 0.019	0.795 ± 0.022	0.849 ± 0.053	0.809 ± 0.014
covtype	0.746 ± 0.005	**0.766** ± 0.003	0.672 ± 0.018	0.674 ± 0.021	0.685 ± 0.016	0.709 ± 0.008
magic04	0.769 ± 0.011	**0.771** ± 0.006	0.734 ± 0.007	0.731 ± 0.015	0.736 ± 0.013	0.752 ± 0.008
heart	**0.883** ± 0.032	0.875 ± 0.026	0.716 ± 0.021	0.753 ± 0.038	0.777 ± 0.043	0.772 ± 0.053
australian	0.841 ± 0.023	**0.842** ± 0.022	0.711 ± 0.056	0.725 ± 0.070	0.742 ± 0.064	0.768 ± 0.036
diabetes	**0.714** ± 0.029	0.705 ± 0.032	0.683 ± 0.037	0.692 ± 0.040	0.694 ± 0.044	0.689 ± 0.040
acoustic	0. 844 ± 0.005	**0.850** ± 0.003	0.840 ± 0.005	0.839 ± 0.002	0.832 ± 0.005	0.841 ± 0.003
vehicle	**0.816** ± 0.018	0.814 ± 0.018	0.764 ± 0.027	0.797 ± 0.014	0.790 ± 0.029	0.805 ± 0.021
segment	**0.838** ± 0.015	0.836 ± 0.008	0.691 ± 0.031	0.768 ± 0.027	0.755 ± 0.024	0.796 ± 0.014

4.4 Results on High-Dimensional Datasets

We study the performance of the proposed CBR-diag$_{FIFO}$ and compare it with online Uni-Exp, OPAUCr, and OAM_{seq} that avoid constructing the full covariance matrix. Table 5 compares our method and the other online algorithms in terms of AUC, while Table 6 shows the classification accuracy at OPTROC. Figure 2 displays the running time (in milliseconds) comparison.

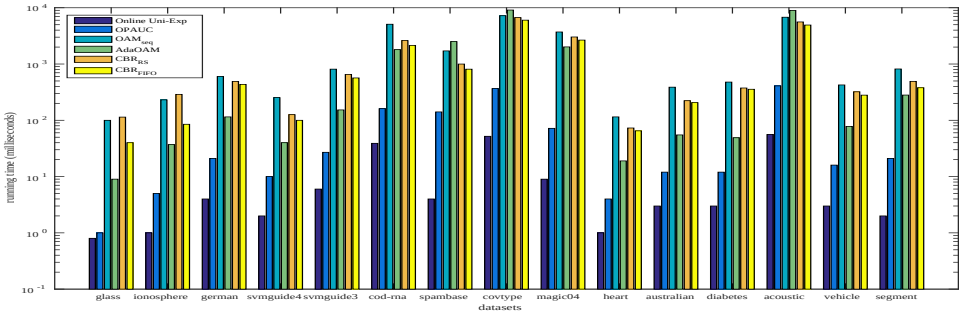

Fig. 1. Running time (in milliseconds) of CBR and the other online learning algorithms on the benchmark datasets. The y-axis is displayed in log- scale.

Table 5. Comparison of AUC on high-dimensional datasets

Data	CBR-diag$_{\text{FIFO}}$	Online Uni-Exp	OPAUCr	OAM$_{\text{seq}}$
farm-ads	**0.961** ± 0.004	0.942 ± 0.006	0.951 ± 0.004	0.952 ± 0.005
rcv1	**0.950** ± 0.007	0.927 ± 0.015	0.914 ± 0.016	0.945 ± 0.008
sector	**0.927** ± 0.009	0.846 ± 0.019	0.908 ± 0.013	0.857 ± 0.008
real-sim	**0.982** ± 0.001	0.969 ± 0.003	0.975 ± 0.002	0.977 ± 0.001
news20	**0.956** ± 0.003	0.939 ± 0.005	0.942 ± 0.006	0.944 ± 0.005
Reuters	**0.993** ± 0.001	0.985 ± 0.003	0.988 ± 0.002	0.989 ± 0.003

The results show that the proposed method CBR-diag$_{FIFO}$ yields a better performance on both measures. We observe that the CBR-diag$_{FIFO}$ presents a competitive running time compared to its counterpart OAM$_{seq}$ as shown in Figure 2. We can also see that the CBR-diag$_{FIFO}$ takes more running time compared to the OPAUCr. However, the CBR-diag$_{FIFO}$ achieves better AUC and classification accuracy compared to the OPAUCr. The online Uni-Exp algorithm requires the least running time, but it presents lower AUC and classification accuracy compared to our method.

Table 6. Comparison of classification accuracy at OPTROC on high-dimensional datasets

Data	CBR-diag$_{\text{FIFO}}$	Online Uni-Exp	OPAUCr	OAM$_{\text{seq}}$
farm-ads	**0.897** ± 0.007	0.872 ± 0.012	0.885 ± 0.008	0.882 ± 0.007
rcv1	**0.971** ± 0.001	0.967 ± 0.002	0.966 ± 0.003	0.970 ± 0.001
sector	**0.850** ± 0.012	0.772 ± 0.011	0.831 ± 0.015	0.776 ± 0.008
real-sim	**0.939** ± 0.003	0.913 ± 0.005	0.926 ± 0.002	0.929 ± 0.001
news20	**0.918** ± 0.005	0.895 ± 0.005	0.902 ± 0.009	0.907 ± 0.006
Reuters	**0.971** ± 0.004	0.953 ± 0.006	0.961 ± 0.006	0.961 ± 0.006

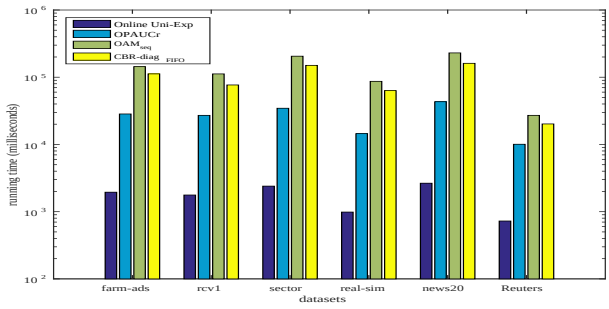

Fig. 2. Running time (in milliseconds) of CBR-diag$_{FIFO}$ algorithm and the other online learning algorithms on the high-dimensional datasets. The y-axis is dis- played in log-scale.

5 Conclusions and Future Work

In this paper, we proposed a linear online soft confidence-weighted bipartite ranking algorithm that maximizes the AUC metric via optimizing a pairwise loss function. The complexity of the pairwise loss function is mitigated in our algorithm by employing a finite buffer that is updated using one of the stream oblivious policies. We also develop a diagonal variation of the proposed confidence-weighted bipartite ranking algorithm to deal with high-dimensional data by maintaining only the diagonal elements of the covariance matrix instead of the full covariance matrix. The experimental results on several benchmark and high-dimensional datasets show that our algorithms yield a robust performance. The results also show that the proposed algorithms outperform the first and second-order AUC maximization methods on most of the datasets. As future work, we plan to conduct a theoretical analysis of the proposed method. We also aim to investigate the use of online feature selection [36] within our proposed framework to effectively handle high-dimensional data.

References

1. Agarwal, S.: A study of the bipartite ranking problem in machine learning (2005)

2. Cai, D., He, X., Han, J.: Locally consistent concept factorization for document clustering. IEEE Transactions on Knowledge and Data Engineering 23(6), 902–913 (2011)
3. Cavallanti, G., Cesa-Bianchi, N., Gentile, C.: Tracking the best hyperplane with a simple budget perceptron. Machine Learning 69(2-3), 143–167 (2007)
4. Cesa-Bianchi, N., Conconi, A., Gentile, C.: A second-order perceptron algorithm. SIAM Journal on Computing 34(3), 640–668 (2005)
5. Chapelle, O., Keerthi, S.S.: Efficient algorithms for ranking with svms. Information Retrieval 13(3), 201–215 (2010)
6. Cortes, C., Mohri, M.: Auc optimization vs. error rate minimization. Advances in neural information processing systems 16(16), 313–320 (2004)
7. Crammer, K., Dekel, O., Keshet, J., Shalev-Shwartz, S., Singer, Y.: Online passive-aggressive algorithms. The Journal of Machine Learning Research 7, 551–585 (2006)
8. Crammer, K., Dredze, M., Kulesza, A.: Multi-class confidence weighted algorithms. In: Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2-Volume 2. pp. 496–504. Association for Computational Linguistics (2009)
9. Crammer, K., Dredze, M., Pereira, F.: Confidence-weighted linear classification for text categorization. The Journal of Machine Learning Research 13(1), 1891–1926 (2012)
10. Crammer, K., Kulesza, A., Dredze, M.: Adaptive regularization of weight vectors. In: Advances in neural information processing systems. pp. 414–422 (2009)
11. Dekel, O., Shalev-Shwartz, S., Singer, Y.: The forgetron: A kernel-based perceptron on a budget. SIAM Journal on Computing 37(5), 1342–1372 (2008)
12. Ding, Y., Zhao, P., Hoi, S.C., Ong, Y.S.: An adaptive gradient method for online auc maximization. In: AAAI. pp. 2568–2574 (2015)
13. Dredze, M., Crammer, K., Pereira, F.: Confidence-weighted linear classification. In: Proceedings of the 25th international conference on Machine learning. pp. 264–271. ACM (2008)
14. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. The Journal of Machine Learning Research 12, 2121–2159 (2011)
15. Gao, W., Jin, R., Zhu, S., Zhou, Z.H.: One-pass auc optimization. In: ICML (3). pp. 906–914 (2013)
16. Hanley, J.A., McNeil, B.J.: The meaning and use of the area under a receiver operating characteristic (roc) curve. Radiology 143(1), 29–36 (1982)
17. Hu, J., Yang, H., King, I., Lyu, M.R., So, A.M.C.: Kernelized online imbalanced learning with fixed budgets. In: AAAI. pp. 2666–2672 (2015)
18. Joachims, T.: A support vector method for multivariate performance measures. In: Proceedings of the 22nd international conference on Machine learning. pp. 377–384. ACM (2005)
19. Joachims, T.: Training linear svms in linear time. In: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 217–226. ACM (2006)
20. Kar, P., Sriperumbudur, B.K., Jain, P., Karnick, H.: On the generalization ability of online learning algorithms for pairwise loss functions. In: ICML (3). pp. 441–449 (2013)
21. Kotlowski, W., Dembczynski, K.J., Huellermeier, E.: Bipartite ranking through minimization of univariate loss. In: Proceedings of the 28th International Conference on Machine Learning (ICML-11). pp. 1113–1120 (2011)

22. Kuo, T.M., Lee, C.P., Lin, C.J.: Large-scale kernel ranksvm. In: SDM. pp. 812–820. SIAM (2014)
23. Lee, C.P., Lin, C.b.: Large-scale linear ranksvm. Neural computation 26(4), 781–817 (2014)
24. Li, N., Jin, R., Zhou, Z.H.: Top rank optimization in linear time. In: Advances in Neural Information Processing Systems. pp. 1502–1510 (2014)
25. Liu, T.Y.: Learning to rank for information retrieval. Foundations and Trends in Information Retrieval 3(3), 225–331 (2009)
26. Ma, J., Kulesza, A., Dredze, M., Crammer, K., Saul, L.K., Pereira, F.: Exploiting feature covariance in high-dimensional online learning. In: International Conference on Artificial Intelligence and Statistics. pp. 493–500 (2010)
27. Orabona, F., Crammer, K.: New adaptive algorithms for online classification. In: Advances in neural information processing systems. pp. 1840–1848 (2010)
28. Orabona, F., Keshet, J., Caputo, B.: Bounded kernel-based online learning. The Journal of Machine Learning Research 10, 2643–2666 (2009)
29. Rendle, S., Balby Marinho, L., Nanopoulos, A., Schmidt-Thieme, L.: Learning optimal ranking with tensor factorization for tag recommendation. In: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 727–736. ACM (2009)
30. Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organization in the brain. Psychological review 65(6), 386 (1958)
31. Sculley, D.: Large scale learning to rank. In: NIPS Workshop on Advances in Ranking. pp. 1–6 (2009)
32. Vitter, J.S.: Random sampling with a reservoir. ACM Transactions on Mathematical Software (TOMS) 11(1), 37–57 (1985)
33. Wan, J., Wu, P., Hoi, S.C., Zhao, P., Gao, X., Wang, D., Zhang, Y., Li, J.: Online learning to rank for content-based image retrieval. In: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI. pp. 2284–2290 (2015)
34. Wang, J., Wan, J., Zhang, Y., Hoi, S.C.: Solar: Scalable online learning algorithms for ranking
35. Wang, J., Zhao, P., Hoi, S.C.: Exact soft confidence-weighted learning. In: Proceedings of the 29th International Conference on Machine Learning (ICML-12). pp. 121–128 (2012)
36. Wang, J., Zhao, P., Hoi, S.C., Jin, R.: Online feature selection and its applications. IEEE Transactions on Knowledge and Data Engineering 26(3), 698–710 (2014)
37. Zhao, P., Jin, R., Yang, T., Hoi, S.C.: Online auc maximization. In: Proceedings of the 28th International Conference on Machine Learning (ICML-11). pp. 233–240 (2011)
38. Zinkevich, M.: Online convex programming and generalized infinitesimal gradient ascent (2003)